

CLAIMS

- 1 1. A method for a storage operating system implemented in a storage system to op-
2 timize the amount of readahead data retrieved for a read stream established in a data
3 container stored in the storage system, the method comprising:
4 receiving a client read request at the storage system, the client read request indi-
5 cating client-requested data for the storage operating system to retrieve from the data
6 container containing the read stream;
7 determining whether the storage operating system is permitted to retrieve reada-
8 head data from the data container in response to the received client read request;
9 if it is determined that the storage operating system is permitted to retrieve reada-
10 head data from the data container, performing the steps of:
11 (i) selecting an amount of readahead data to retrieve from the data container
12 based on one or more factors; and
13 (ii) retrieving the selected amount of readahead data from the data container.
- 1 2. The method of claim 1, wherein the data container is a file, directory, vdisk or
2 lun.
- 1 3. The method of claim 1, wherein the storage operating system is determined to be
2 permitted to retrieve readahead data from the data container when the client-requested
3 data extends the read stream past a predetermined next readahead value.
- 1 4. The method of claim 3, wherein the predetermined next readahead value is stored
2 in a readset data structure associated with the read stream.
- 1 5. The method of claim 3, wherein the predetermined next readahead value is up-
2 dated based on a percentage of the selected amount of readahead data.

1 6. The method of claim 1, wherein a read-access style associated with the data con-
2 tainer is one of the one or more factors used to select the amount of readahead data.

1 7. The method of claim 6, wherein the selected amount of readahead data equals
2 zero if the read-access style corresponds to a random read-access style.

1 8. The method of claim 1, wherein a number of client read requests processed in the
2 read stream is one of the one or more factors used to select the amount of readahead data.

1 9. The method of claim 8, wherein the number of client read requests processed in
2 the read stream is stored as a count value in a readset data structure associated with the
3 read stream.

1 10. The method of claim 1, wherein the amount of client-requested data is one of the
2 one or more factors used to select the amount of readahead data.

1 11. The method of claim 10, wherein the selected amount of readahead data is set
2 equal to a predetermined upper limit for large amounts of client-requested data.

1 12. The method of claim 1, wherein the selected amount of readahead data is doubled
2 if the number of client read requests processed in the read stream is greater than a first
3 threshold value.

1 13. The method of claim 1, wherein the client-requested data is identified as read-
2 once data when either (i) the number of client read requests processed in the read stream
3 is greater than a second threshold value or (ii) a set of metadata associated with the read
4 stream indicates that the client-requested data is read-once data.

1 14. The method of claim 1, wherein the selected amount of readahead data is stored in
2 one or more buffers enqueued on a flush queue, the flush queue being configured to reuse
3 buffers after a predetermined period of time.

1 15. The method of claim 14, wherein the predetermined period of time equals two
2 seconds.

1 16. An apparatus configured to implement a storage operating system that optimizes
2 the amount of readahead data retrieved for a read stream established in a data container
3 stored in the apparatus, the apparatus comprising:

4 means for receiving a client read request, the client read request indicating client-
5 requested data for the storage operating system to retrieve from the data container con-
6 taining the read stream;

7 means for determining whether the storage operating system is permitted to re-
8 trieve readahead data from the data container in response to the received client read re-
9 quest;

10 means for selecting an amount of readahead data to retrieve from the data con-
11 tainer based on one or more factors; and

12 means for retrieving the selected amount of readahead data from the data con-
13 tainer.

1 17. The apparatus of claim 16, wherein the data container is a file, directory, vdisk or
2 lun.

1 18. The apparatus of claim 16, wherein the storage operating system is determined to
2 be permitted to retrieve readahead data from the data container when the client-requested
3 data extends the read stream past a predetermined next readahead value.

1 19. The apparatus of claim 18, further comprising means for updating the predeter-
2 mined next readahead value based on a percentage of the selected amount of readahead
3 data.

1 20. The apparatus of claim 16, wherein the one or more factors used to select the
2 amount of readahead data includes at least one of:

- 3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 21. The apparatus of claim 16, wherein the selected amount of readahead data is dou-
2 bled if the number of client read requests processed in the read stream is greater than a
3 first threshold value.

1 22. A storage system configured to optimize the amount of readahead data retrieved
2 for a read stream established in a data container stored in the storage system, the storage
3 system comprising:

4 a network adapter for receiving a client read request, the client read request indi-
5 cating client-requested data to retrieve from the data container containing the read stream;
6 and

7 a memory configured to store instructions for implementing a storage operating
8 system that performs the steps of:

9 determining whether the storage operating system is permitted to retrieve
10 readahead data from the data container in response to the received client read re-
11 quest, and

12 if it is determined that the storage operating system is permitted to retrieve
13 readahead data from the data container:

- 14 (i) selecting an amount of readahead data to retrieve from the
15 data container based on one or more factors; and

16 (ii) retrieving the selected amount of readahead data from the
17 data container.

1 23. The storage system of claim 22, wherein the data container is a file, directory,
2 vdisk or lun.

1 24. The storage system of claim 22, wherein the storage operating system is deter-
2 mined to be permitted to retrieve readahead data from the data container when the client-
3 requested data extends the read stream past a predetermined next readahead value.

1 25. The storage system of claim 24, wherein the predetermined next readahead value
2 is updated based on a percentage of the selected amount of readahead data.

1 26. The storage system of claim 22, wherein the one or more factors used to select the
2 amount of readahead data includes at least one of:

- 3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 27. The storage system of claim 22, wherein the selected amount of readahead data is
2 doubled if the number of client read requests processed in the read stream is greater than
3 a first threshold value.

1 28. A computer-readable media comprising instructions for execution in a processor
2 for the practice of a method for a storage operating system implemented in a storage
3 system to optimize the amount of readahead data retrieved for a read stream established
4 in a data container stored in the storage system, the method comprising:
5 receiving a client read request at the storage system, the client read request indi-
6 cating client-requested data for the storage operating system to retrieve from the data
7 container containing the read stream;

8 determining whether the storage operating system is permitted to retrieve reada-
9 head data from the data container in response to the received client read request;
10 if it is determined that the storage operating system is permitted to retrieve reada-
11 head data from the data container, performing the steps of:
12 (i) selecting an amount of readahead data to retrieve from the data container
13 based on one or more factors; and
14 (ii) retrieving the selected amount of readahead data from the data container.

1 29. The computer-readable media of claim 28, wherein the data container is a file,
2 directory, vdisk or lun.